



LIGO Data & Computing Update

Albert Lazzarini
LIGO Laboratory at Caltech

12th Meeting of the LIGO Laboratory PAC
26 June 2002
Cambridge, MA

- Simulation & Modeling
- Data Analysis Systems
- Grid Computing
- Staffing



LIGO Modeling & Simulation Activities Overview

- e2e development began after LIGO I design was essentially complete
 - » e2e used by M. Evans in successful redesign of LIGO I lock acquisition control system
- Major ongoing effort at present:
 - » Model refinement, speed-up of code
 - » Representation of realistic noise performance of the locked state interferometer
- For the future:
 - » Integration of CDS real time code for *as-built* LIGO length and alignment control system into the simulated interferometer model
 - Discussion has just begun ...
 - But not likely to happen due to limited time, manpower resources

First generation LIGO simulation: *Han2k*

- Purpose
 - » Design and develop the LHO 2 km interferometer locking servo
 - » Simulate major characteristics of length degree of freedom below 20 Hz.
- Simulation includes
 - » Scalar field approximation
 - 1 DOF, TEM₀₀ throughout model
 - » Saturation of actuator drivers
 - » Simplified seismic motion with low-f correlations
 - » Analog length control, with no angular control (i.e., 1 DOF)
 - » No additional noise sources
 - no frequency noise, shot noise, sensor/actuator/electronic noise



LIGO Second generation LIGO simulation: *SimLIGO*

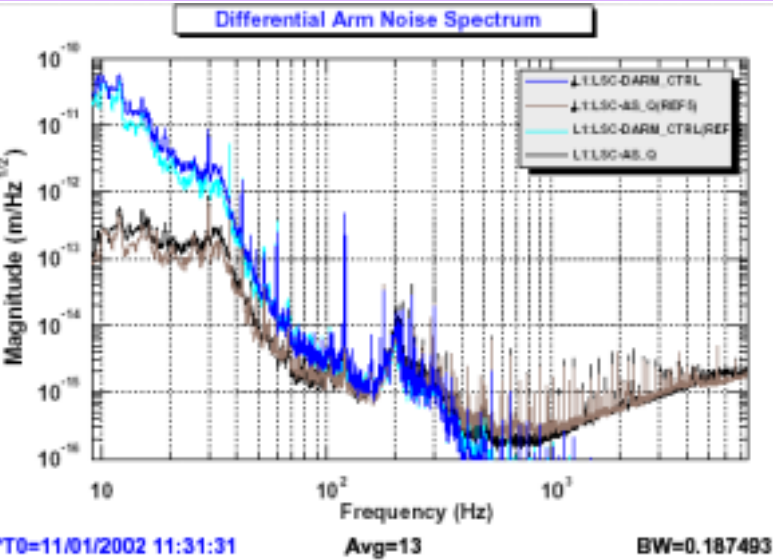
- Purpose
 - » Quantitative noise and performance estimation of as-built interferometer
 - » Assist noise identification and reduction, lock stability studies during commissioning
- Simulation includes
 - » As-built optical and mechanical parameters
 - » Seismic motion correlations among chambers
 - » 3D mirror with 4 sensors and actuators
 - » Digital length control system
 - » Alignment control via wavefront sensors and/or optical lever
 - » Digital suspension controller
 - » Common mode servo
 - » Mode cleaner with suspended, controlled small optics

- Interferometric
 - » Optical asymmetries
 - reflectivities, transmissivity, length, radii of curvature, phase maps, ...
 - Non-normal incidence -- wedge angles, Earth curvature
 - » Scattered light

- Mechanical
 - » Wire resonances, test mass internal modes

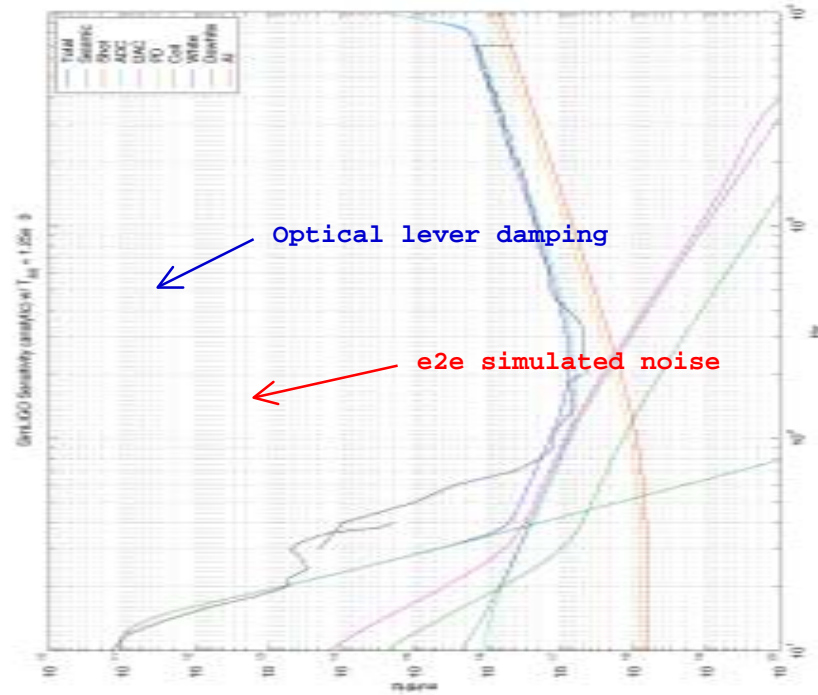
- Sensing, actuation, digital real-time electronics
 - » Photo-detector, coil drivers
 - » Whitening/de-whitening filters, anti-aliasing
 - » ADC, DAC, digital transfer functions

Simulation of sensitivity curve noise by realistic simulation



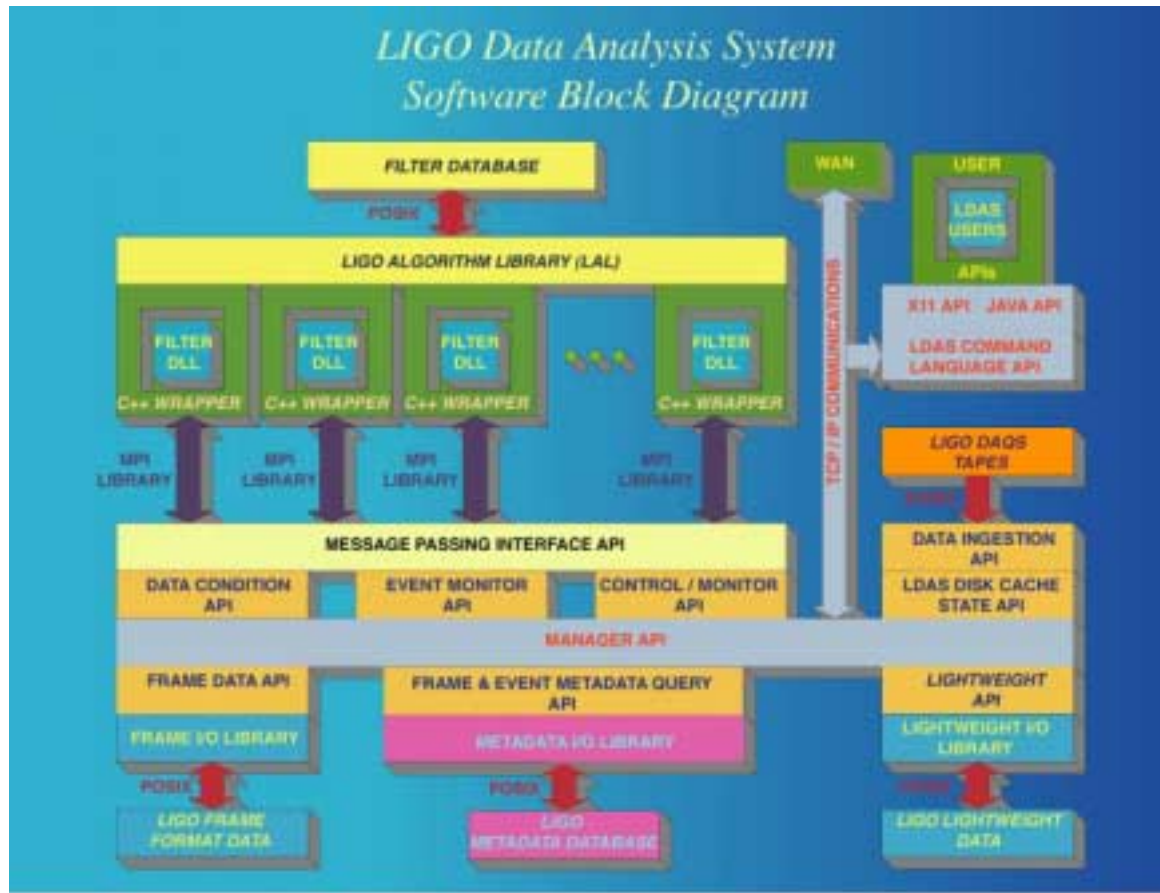
Simulation includes:

- Locked interferometer
- Mechanics
- Sensor-actuator, optical lever damping
- Servo electronics
- Signal extraction
- Simulated and measured noises



- Evaluation of the as-built LIGO performance
 - » Identify minimum achievable noise performance for the as-built design
 - » Study, quantify specific noise sources and issues:
 - bi-linear couplings, non-linearities, non stationarity, transients, lock instabilities, etc
 - » LIGO simulation package is ready
 - ✓ Length control: complete
 - Alignment control: available soon
- Ultimate success determined by adoption & use ...
 - » Hardware experts know the important problems
 - Lock acquisition was successful because it was a collaboration between simulation and interferometry experts
 - Simulation group is small and needs to leverage off other within laboratory, LSC
 - » Documentation is almost complete
 - How-to pages: setting up new models, running, ...
 - Tutorial, SimLIGO manual, physics models embedded in e2e simulation, ...

Data Analysis - Software -

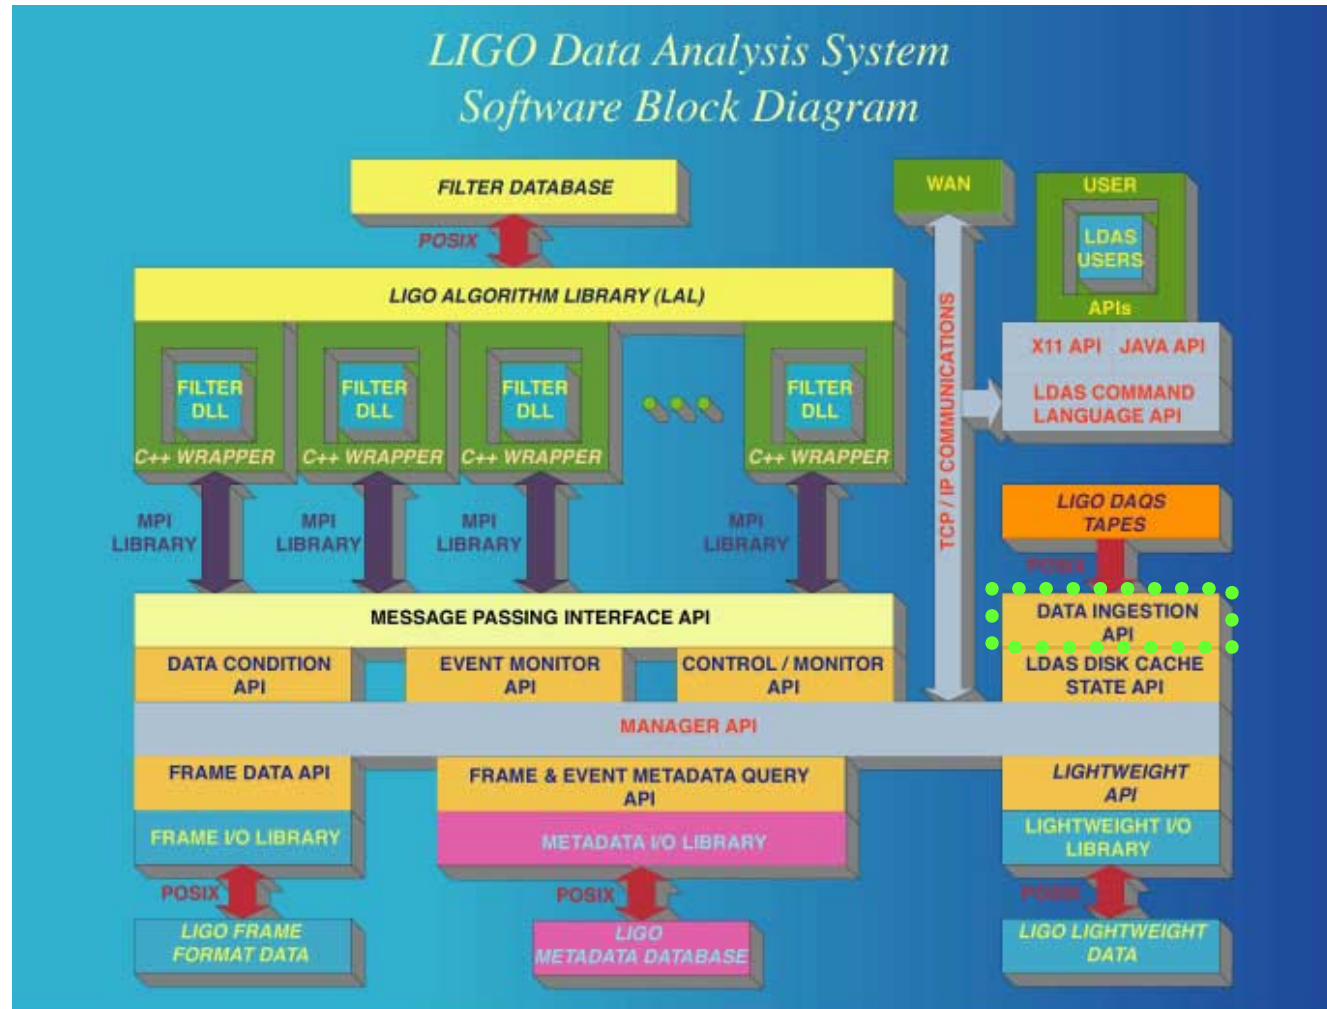


LDAS-0.3.0

- Only one module left to build

dataIngestionAPI:

- Basic functions needed with the beginning of the Engineering Runs.
- Scripts external to LDAS developed to carry out these functions.
- Script to recording data to tapes at sites can be controlled and monitored using LDAS's controlMonitorAPI.



E7 LDAS Job Summary

28 December 2001 - 14 January 2002

	Hanford LDAS	Livingston LDAS	MIT LDAS	CIT-TEST LDAS	TOTAL
Total Jobs	63600	48775	280	915	113570
Database Rows	4188188	2789132	1062	2096	6980478

- **LDAS version 0.0.23 (alpha code) ran for full E7 Run:**
 - » 28 December 2001 - 14 January 2002
 - » Approximately **1 job every 10 seconds** (averaged).
 - » Approximately **50 DB rows every 10 seconds** (averaged).
- **> 90% of jobs completed successfully**
 - » LHO: ~ 92%; LLO: ~95%; Not checked elsewhere.
- **However ... pre-release testing predicted ~0.3% failure rate ...**
 - » Dominated by dataConditionAPI thread problems.
 - » A fraction also due to mpiAPI/wrapperAPI communications issues.
 - » Rare intermittent issues (difficult to debug) also suspected.

Reduced Data Set (RDS) Production

7 - 9 June 2002

- Ran with LDAS version 0.2.0 (first beta release) during E8
- University of Oregon (Isabel Leonor) led effort to write RDS frames
 - » Approximately 120 channels out of the 5000 selected out.
 - » Approximately 20 fast channels re-sampled by factor of two.
 - » Each RDS frame 20% smaller than corresponding raw frame.
- LDAS 0.2.0 unable to keep up with data rate as it wrote RDS frames!
- Repeated Test 2 weeks later at Hanford using pre-release version 0.2.55 of LDAS
 - » Increased number of channels to 152 (per S1 Run Plan)
 - » LDAS able to process 80 seconds data in 35 seconds when IFO in lock ... however ...
Discovered that **out of lock** processing time grew from 35 seconds to 155 seconds
 - Isolated to sample values below 10^{-24} in the DARM control channel tickling a peculiarity in floating point performance on Suns using GCC compiler.
 - Value of 10^{-24} in DARM written by realtime control system to replace “0” in order to avoid 1/0 in matrix inversion.

- LDAS at MIT has become central to analysis activities within the BURST Upper Limits Group.
 - » See http://www.ligo.caltech.edu/LIGO_web/0203news/0203mit.html
- *LDAS used by Sylvestre at MIT to analyze E7 data and contributed to his recent Ph. D. dissertation.*
- Burst Group worked with LDAS development team to test pre-Science Run version of LDAS.
- As part of the analysis of E7 data, 6.5 days of E7 data were analyzed in a single day:
 - » 1554 jobs submitted to LDAS-MIT and completed successfully
 - » (3 DSOs) x (2 IFOs) x (1 waveform) x (7 amplitudes) x (37 360 sec intervals)



- ❑ Individual accounts provided to individual users
 - No more shared user accounts!
 - Addresses long standing issue raised in NSF review.

File

Select controlmon server at:

LDAS control | Status | System Load | Jobs | LDAS Logs | MPI Jobs | Beowulf System | Tape Control

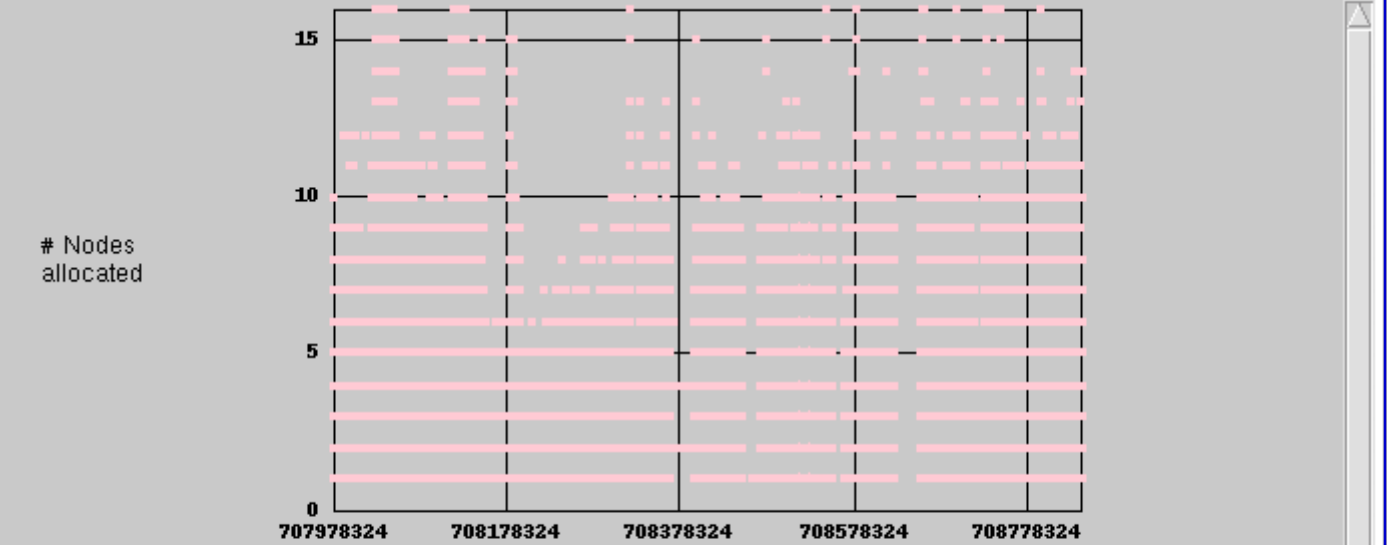
LDAS APIs | LDAS system | LDAS Resources | LDAS utilities | User Admin | LDAS Core Files

80740 Insertions, 7094 Queries, for all table(s) for user(s) all from 708755574-708841974

% Failed: 0.08, 11776 Jobs, 9 failed for user(s) all for times 708755432-708841832

API Legends: manager diskcache frame metadata ligolv ligo datacond eventmon cntlmon

From 06/12/02 21:38:31 PDT-06/22/02 21:38:31 PDT, 48325 jobs



Close Save zoom out zoom in


Control and Monitor API

- user accounts management
- DB2 usage statistics
- user job statistics
- API memory usage
- MPI node usage statistics

- Third beta release of LDAS will be used (version 0.3.0)
- Exhaustive testing of release for three weeks prior to S1:
 - » First week consisted of internal testing and bug fixing
 - » Second week continued internal testing and began integration testing with LAL and LALwrapper codes
 - » Carried out collaborative RDS testing at LHO and Burst Group analysis testing at MIT in the third week.
 - » Pushed pre-release 0.2.58 to UWM one week before S1 Run.
- Push release to all Laboratory sites (just) before S1 begins.
 - » “Just in time delivery”
- LSC (through LSUG group) will work with LDAS to “ramp” up to the S1 Run analysis plan.

- Version 5 Frame Specification significantly delayed and has resulted in some loss of schedule and reduced functionality for the Science Run.
- Reduced Data Sets are not very reduced!
 - » Frame size (>20% of raw data) taxing the current capabilities of LDAS and will demand reworking some of the frameAPI to be able to support this load.
- It is demanding to develop beta level code while supporting use of LDAS for science *and* supporting growing population of LSC users, LDAS installations.
- LSC participation in code (algorithm) development has significantly decreased due to change of focus towards Upper Limits Group activities.
 - » Needed signal processing functionality not yet unimplemented: Kalman filter, SI units management, process history, filter coefficient tracking, test signal generation, data type conversion (real->int), data alignment - handling filter delays
- LDAS use model evolution continues
 - » continued need for programming activities vis-à-vis science

Data Analysis - Hardware -



LIGO-G020266-0

“Increase computational capacity over E7 and investigate advanced storage configurations but delay full compute farm deployment until S3.”

	SAN (TB)	IDE (TB)	CPU (Aggregate GHz)	Tape (TB)
LHO	10		139	2
LLO	5		107	2
CIT	3		34	90
MIT	1	2	45	0
DEV	1		25	2
TEST	1		8	0

SAM-QFS vs HPSS

Last major trade study, to be completed by S2:

- SAM-QFS - Sun Microsystems acquisition (LSC), provides alternative to HPSS for integrated mass storage (disk+tape)
 - » http://www.sun.com/products-n-solutions/hardware/docs/Software/Storage_Software/LSC/index.html
- SAM-QFS advantages
 - » Simplicity/reliability
 - » Media import/export
 - » License cost allows for use at observatories
 - » Disaster recovery (GNU TAR)
 - » Metadata performance (x1000)
 - » Single vendor solution (server, software and OEM storage)
 - » Reduced dependency on CACR
- HPSS advantages
 - » Integrated several man-years experience
 - ...which is why we are looking at SAM-QFS ...
 - » Free at Caltech
 - cannot be used other Laboratory sites
 - » **35TB** successfully stored to date

- Additional Fibre Channel ports to connect remaining 2TB of disk at LHO and 1TB at LLO.
 - Enables direct access to the raw frames from additional servers, e.g., CDS and GC.
- Tape drives for archive system at Caltech.
 - The present prototyping set of 2 STK9940A drives will be enhanced with ~8 STK9940B drives in time for S2.
- Tape robotics at LHO and LLO.
 - If SAM-QFS is chosen over HPSS the same tape drive technology will be installed at the Observatories as at the central archive at Caltech.
 - Allows the direct ingestion of tapes without having the need for media-to-media copying.
- Large volume RAID storage @ Caltech, in front of tape archive.
 - Use inexpensive disk storage (\$4k/TB) to build up ~50 - 100 TB of spinning storage in front of tape archive
 - Consider this “consummable” -- build up as needed as data volume grows.
- S3 Beowulf clusters :
 - The initial LDAS science run hardware configuration will be purchased in 4Q02-1Q03 and comprise ~400 compute nodes.
 - The interesting choice between ~3GHz P4 and 64-bit AMD/Intel platforms (if available).

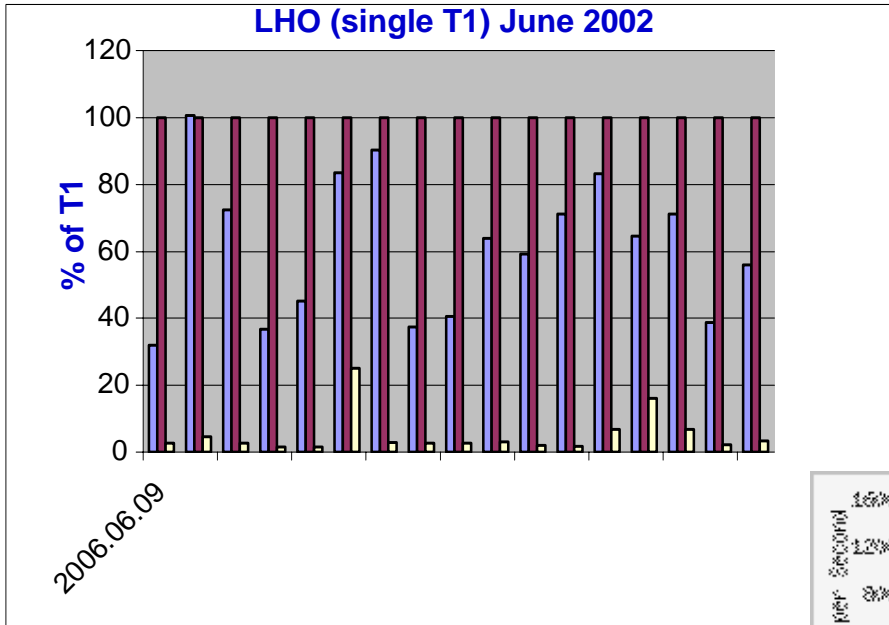
- S3 Server configuration:
 - Data conditioning servers will be upgraded to Science Run configuration.
 - The desire to get to 64-bit platforms may push this out until 2-3Q03 (and hence expend monies from Operations rather than Construction -- *treat as an upgrade*).
 - Accrued experience with Sun V880 SPARC servers shows full more powerful configurations needed
 - Currently 50% populated with 4 processors
 - Additional 900-1000MHz CPUs will be added to those servers that have a demonstrated need, i.e., database, gateway, and/or dataserver.

- Inter-LDAS communications:
 - Need to facilitate the direct communication between database instances at different LDAS sites
 - Support efficient DB federation and replication
 - Laboratory-wide LDAS Virtual Private Network (VPN) will be installed.
 - Initial tests have been with OpenBSD on Intel hardware but this may change.


Data Analysis

- Wide Area Network (WAN) -

- Both LLO, LHO site connections are becoming saturated:
 - » T1 at LHO provided under MOU with DOE/ESnet
 - » 2xT1 at LLO provided through State of Louisiana Dept. of Telecommunications, via LSU.
- Increased traffic:
 - » Collaboration access to on site data
 - » Video for remote control rooms at Caltech, MIT
 - » LDAS data transmission, SW synchronization
- FY 2002 - 2006 Operations proposal contained budget for WAN upgrade to OC3
 - » Recommended also by previous NSF review panels
- Budget item *(for entire FY2002 - 2006 period)* eliminated to meet reduced funding profile guidance from NSF
- Presently exploring options for bandwidth upgrades that can fit into current budgetary constraints
 - » Identify non-recurring & recurring costs, develop a Change Control Board request and formulate a compelling argument for re-allocating funds for WAN upgrade



LLO (dual T1), June 2002.



- LHO frequently saturates
- LLO is at 50% of 2xT1 during peak periods

- LHO

- » Landscape of options has changed markedly over past 18 months
- » New PUC enterprise, Energy Northwest (ENW), has installed fibre-optic infrastructure along the power grid rights-of-way throughout Washington
- » Provides good service at reasonable cost for government, non-profit organizations - Battelle/DOE have already migrated to ENW
- » Cost of OC3 is expected to be ~13X less per Mbps than present T1 costs:
 - OC3 annual costs would be additional \$84k/year over what LIGO pays for T1
- » LIGO will continue its MOU with DOE/ESnet and share infrastructure, BW with PNNL/Battelle
 - Mutual cooperation between NSF, DOE deemed important to PNNL, LIGO
- » Awaiting quote for service from PNNL, ENW in order to formulate a proposal to LIGO Laboratory Directorate

- LLO
 - » Unlike LHO, the LA State infrastructure is linked to commercial enterprise
 - Market rate for bandwidth in southeastern US is higher
 - Fewer options (only Bell South)
 - Several expected competitors aborted plans after the *.com* collapse in FY2000-2001
 - » Presently in a “wait and see” mode, while identifying options
 - » Two factors mitigate situation vis-à-vis LHO:
 - Data rate, volume is 0.5 as great
 - Present bandwidth is already twice as great
 - 2 T1 lines in place vs. 1 at LHO

Data Analysis

- Grid Computing -

GriPhyN: GriPhyN: Grid Physics Network

iVDGL: International Virtual Data Grid Laboratory

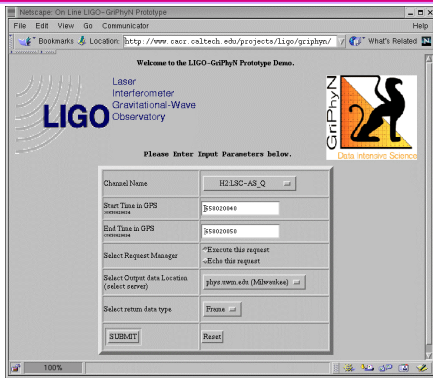
Grid activities

GriPhyN & iVDGL

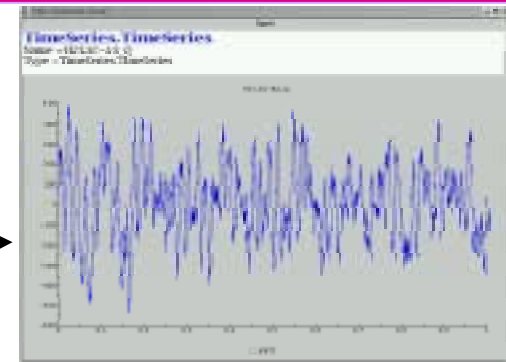
- LIGO/Caltech, UWM, PSU members of collaborations
- GriPhyN - CS + Applications research focused on:
 - » Virtual data in LIGO analysis
 - Use of transformation catalogs to keep track of previously requested, performed transformations that can be used to accelerate data requests
 - Determine if the data product is instantiated, if so where, if not, how to compute it.
 - Plan & execute data movements and computations required to support specific analyses
 - » LIGO data mirroring using gridtools
 - Automated, robust replication between Tier 1 & Tier 2 sites
 - » Developing an interface between the grid environment and the native LDAS environment
 - Extending LDAS into a grid-enabled environment -- longer term goal
 - » Grid security infrastructure
 - Provide data access only to collaboration members
 - Provide a level of security to LDAS computing

GriPhyN Prototype

SC01 Virtual Data Replication Demonstration

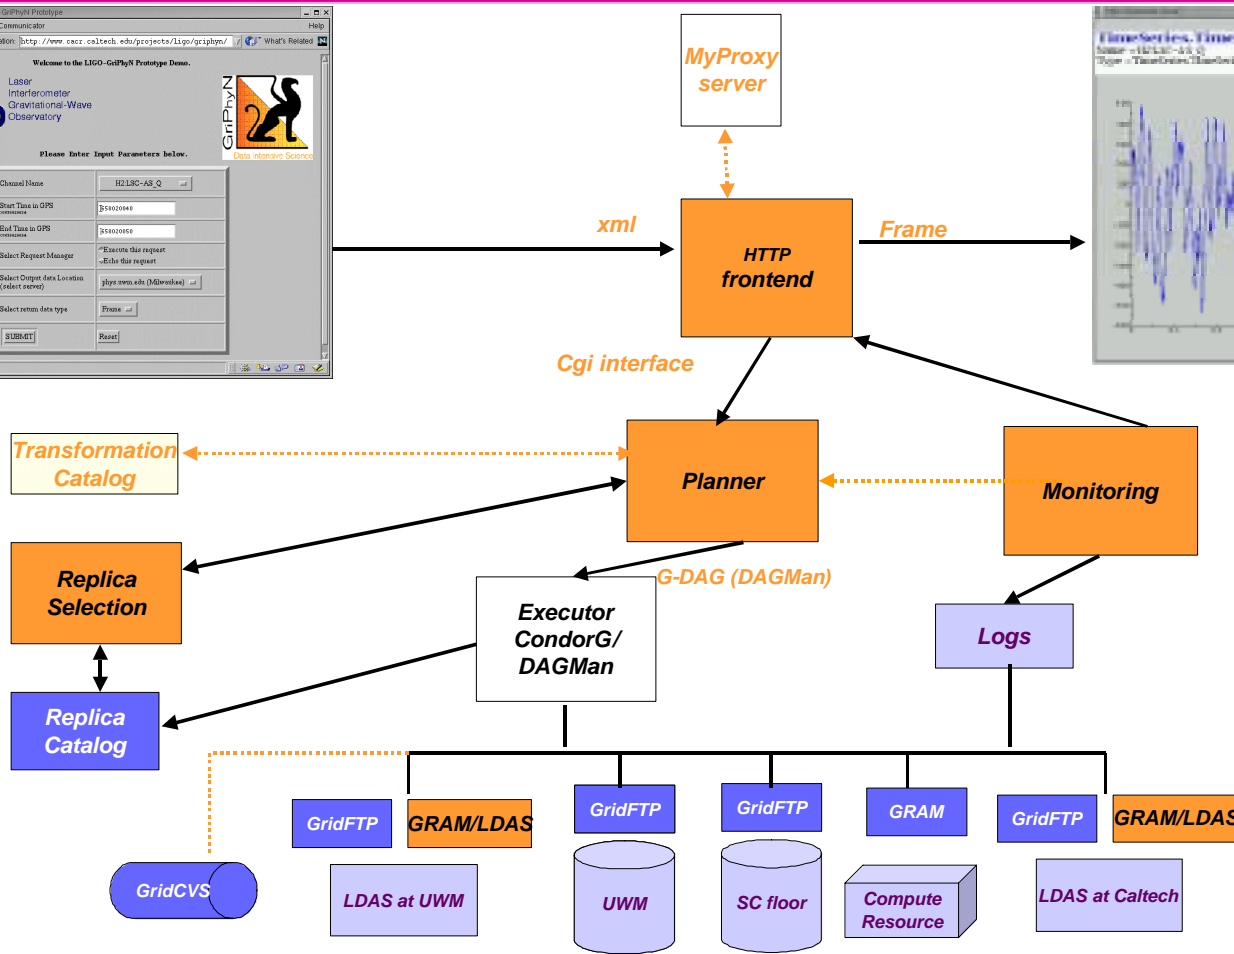


MyProxy server



Desired Result

Single channel time series



Virtual Data Prototype Functionality

- The detailed prototype functionality was as follows:
 - » User inputs request using a web browser. User can specify:
 - Data channel name
 - Time of interest
 - Desired output data location
 - » Request is transformed into XML
 - » Request Interpreter: Understand an XML-specified request
 - » Acquire user's proxy credentials
 - » Replica Selection: Consult replica catalog for available data, select replica "closest to the desired output location"
 - » Request Planner:
 - If data available: plan necessary data movement
 - Else: Construct a plan to produce data not available, including execution location selection. Select input data location ("close to compute resources"), schedule data movements etc...
 - Specify plan in DAGMan format (a Condor-G specification format)
 - » Request Executor: submit DAGMan specified plan to Condor-G
 - » Return requested data to the user specified location in Frame format
 - » Provide a graphical view of the data by using the XSIL frame viewer.

- iVDGL will provide resources to deploy, support (in a limited fashion) the first 2 LIGO Tier 2 centers:
 - » UWM
 - extension of existing MRI-funded hardware - provides out-year hardware renewal
 - » PSU
 - greenfield Tier 2 center to be implemented by end of CY2002
 - » Limited FTE support to operate centers
 - Postdocs, students
 - Matching university funds from PSU will provide limited IT support

- LIGO and Virgo have agreed to implementation of a data exchange protocol based on grid technology
 - » Exchange limited environmental channels to explore geophysical tele-correlations in anticipation of future network analysis using Virgo, LIGO as an array.
 - » Current prototype system developed by LIGO postdoc, CACR scientist, and Virgo collaborators
- Plan is to work with Virgo to migrate to grid toolkit to provide more robust, automated data exchanges around the clock
 - » Fits into the US-EU grid collaboration strategy

- ITR 2003 announcement of opportunity expected from NSF this summer
- LSC Computing Committee organizing a collaboration-wide proposal to request funding to operate LSC Tier 2 centers for LIGO Science Run(s)
 - » *ITR2000: GriPhyN ->C/S + Applications R&D, prototyping*
 - » *ITR2001: iVDGL-> Center buildup (2 for LSC), port, install GriPhyN deliverables to Tier 2 centers*
 - » ITR2003: request manpower to operate centers **to do the science**
 - Scientists, systems administration (~ 10 FTE across 2 LSC sites, Lab. sites)
 - Help desk, 7x24 operations support

Data & Computing - Personnel -

Data & Computing Personnel

	MIT			LHO				
	Classification	FTE	Classification	FTE	Classification	FTE		
Lazzarini	Sci/Mgr	1	Katsavounou	Faculty	-	Mendell	Sci	1
Blackburn	Sci/TskLdr]	1	Bayer	SysAmdin	0.5	1 Yakushin	Sci	1
Anderson	Sci/TskLdr	0.9				Roddy	SysAmdin	0.5
Shawhan	Sci	1						
Charlton	Sci/Postdoc	1						
Wen	Sci/Postdoc	0.5 (+ 0.5LISA)						
Creighton	Sci/Postdoc	0.5 (+ 0.5LISA)						
Sylvestre	Sci/Postdoc	0.5 (+ 0.5LISA)						
Pan	GRStudent	1 (Prince)						
Hua	GRStudent	1 (Prince)						
Maros	SWEng	1						
Ehrens	SWEng/Contract	1						
Barnes	SWEng/Contract	1						
Lei	SWEng/Contract	1						
Salzman	SWEng/Contract	1						
McWhorter	SWEng/Contract	1						
Wilson	SysAdmin/Contr	1						
Kozak	SysAdmin/Contr	1						
Yamamoto	Sci	1						
Bhawal	Sci	1						
Sears	SWEng	1						
Evans	SWEng	1						
Araya	SWEng/Contract	1						
Wallace	SysAdmin	1	Bayer	SysAmdin	0.5	Patton	SysAmdin	0.5
Bogue	SysAdmin/Contr	1						
Pedraza	SysAdmin/Contr	1						
Kondrashov	WebAdmin	0.8						

LIGO-G020266-00-E

FINIS

Significant Performance Improvements to 0.3.0

- Implemented TCL Channel for data sockets communications.
 - » Corrected almost all FILO (first in, last out) staging of jobs
 - » Tripled average throughput rate of jobs in busy system
- Reduced overhead for MPI from over 35 seconds per job to under 5 seconds.
- Removed known thread safety bugs, increasing concurrency in dataConditionAPI.
- Fixed code in-lining and optimization issues to allow for fully optimized build and running of LDAS.
- Removed major memory leaks resulting in greater duty cycle for individual LDAS APIs.



LDAS computational resources planned to be available for on-line processing during S1

LHO

machine

Gateway
Metaserver
Dataserver

Datacon
Beowulf
Node1-64

services

managerAPI/cmonAPI/WWW/NFS
metadataAPI/DB2
frameAPI/diskcacheAPI
raw frames
reduced frames
data archive
datacondAPI
mpiAPI
search DSO's

hardware

E450 4x400MHz/4MB/4GB/560GB
E450 4x400MHz/4MB/4GB/420GB
V880 4x750MHz/8MB/8GB
8.5TB
2.1TB
30 slot/2xAIT-2 robot
Dell6400 4x700MHz/1MB/2GB
Dell2400 2x933MHz/256kB/1GB
64ea. 1x2GHz/512kB/512MB

LLO

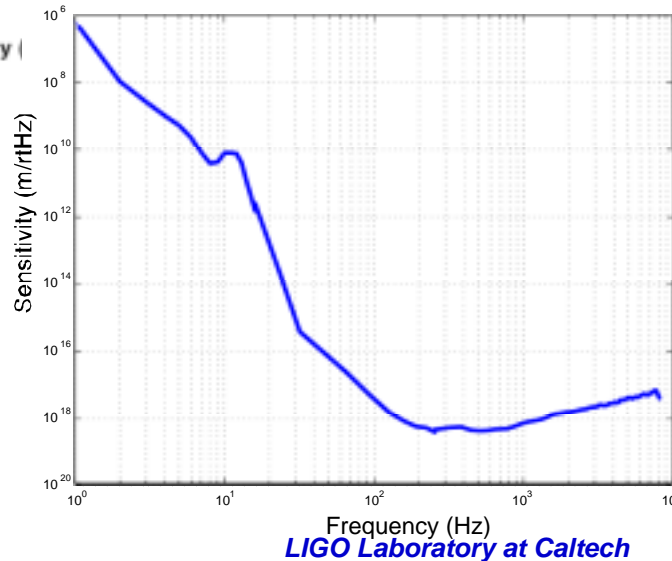
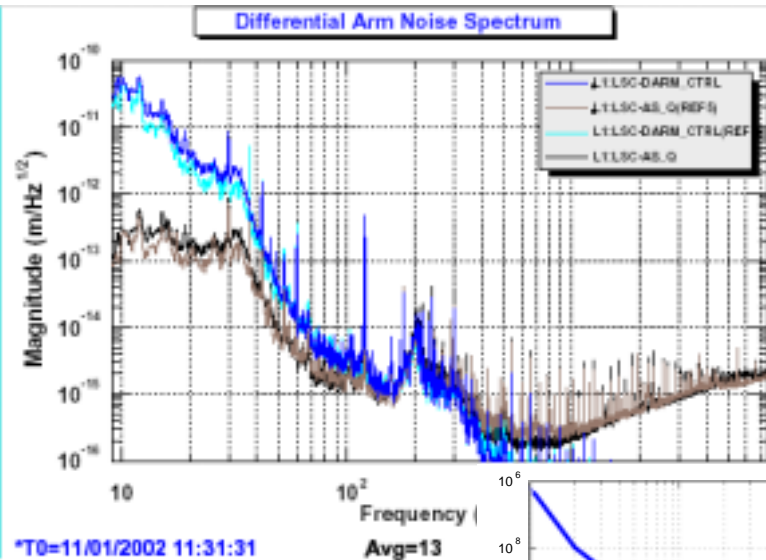
- Same except: (i) only 48 nodes; (ii) 0.5 the dataserver disk storage; (iii) metaserver has 560GB.

Networks

- All servers running Gigabit Ethernet
- All compute nodes running Fast Ethernet
- Dataserver and Gateway running 1Gbit/s Fibre Channel

Simulation of sensitivity curve

noise by realistic simulation



Simulation includes

- Locked interferometer
- Mechanics
- Sensor-actuator
- Servo electronics
- Signal extraction
- Simulated and measured noises
- ...

Simulation of sensitivity curve noise by realistic simulation

